

ABSTRACT

Techniques are disclosed for storing content in distributed computing environments to reduce storage overhead and improve efficiency and/or speed of accessing and delivering the stored content. A content management system (or other authoring system, such as a text editor) supplies hints about dependencies among objects or files, such as which objects are embedded into a web page and which objects are referenced by the web page. These hints are then used when deploying the corresponding files on a storage system, such that the files which are likely to be retrieved together are (preferably) stored together. Optionally, run-time observations by a web server may be used as input to this process to determine how best to store content during a subsequent redeployment.